Reduced storage nodal discontinuous Galerkin methods on semi-structured prismatic meshes

Jesse Chan[a,*], Zheng Wang[a], Russell J. Hewett[c], T. Warburton[b]

[a]*Department of Computational and Applied Mathematics, Rice University, 6100 Main St, Houston, TX, 77005*
[b]*Department of Mathematics, Virginia Tech, 225 Stanger Street Blacksburg, VA, 24061-0123*
[c]*TOTAL E&P Research and Technology USA, 1201 Louisiana St, Houston, TX 77002*

arXiv:1607.03399v2 [math.NA] 1 Nov 2016

Abstract

We present a high order time-domain nodal discontinuous Galerkin method for wave problems on hybrid meshes consisting of both wedge and tetrahedral elements. We allow for vertically mapped wedges which can be deformed along the extruded coordinate, and present a simple method for producing quasi-uniform wedge meshes for layered domains. We show that standard mass lumping techniques result in a loss of energy stability on meshes of vertically mapped wedges, and propose an alternative which is both energy stable and efficient. High order convergence is demonstrated, and comparisons are made with existing low-storage methods on wedges. Finally, the computational performance of the method on Graphics Processing Units is evaluated.

1. Introduction

While time-dependent wave propagation and seismic applications have traditionally utilized high order finite difference solvers, interest in high order finite element methods in seismic imaging has grown over the past two decades, in part due to the ease of adapting finite element methods to complex geometries. However, compared to finite difference methods, a drawback of traditional finite element methods for time-domain wave propagation problems is that a global mass matrix system must be solved at every timestep. Several new high order finite element methods were introduced to address this issue. Among these methods, a commonly used high order method is the spectral element method (SEM) [1, 2], which avoids the solution of a global system by employing a diagonal mass-lumped approximation to the mass matrix. However, the use of SEM is currently limited to unstructured hexahedral meshes, which can be difficult to construct for arbitrary topologies.

Unlike spectral elements methods, high order discontinuous Galerkin (DG) methods [3] can use more general meshes. Like spectral methods, DG methods avoid the inversion of a globally coupled mass matrix. Additionally, due to their low communication costs and high arithmetic intensity, time-domain DG methods have also been shown to be well-suited to parallelization, both on distributed and many-core architectures. Klöckner et al. presented an efficient implementation of quadrature-free nodal DG methods for tetrahedral meshes on a single Graphics Processing Unit (GPU) [4], while the scalability of the multi-GPU case was investigated in [5].

The utility and efficiency of GPU-accelerated nodal discontinuous Galerkin methods for seismic applications on tetrahedral meshes is explored in detail by Modave et al. in [6, 7]. However, the GPU implementation of DG methods on non-tetrahedral elements requires some care, as storage requirements for such elements can be quite large, especially at high orders of approximation. Since many-core accelerators typically have

*Principal Corresponding author
Email addresses: Jesse.Chan@rice.edu (Jesse Chan), zw14@caam.rice.edu (Zheng Wang), russell.hewett@total.com (Russell J. Hewett), tcew@vt.edu (T. Warburton)

a relatively small amount of on-device storage, storage-heavy schemes can greatly limit the maximum problem size which is executable on a single GPU. Furthermore, this can be detrimental to performance, as accelerators such as GPUs typically require sufficiently large problem sizes to operate at peak efficiency.

In [8], a GPU-accelerated time-domain DG method was introduced for hybrid meshes consisting of hexahedra, tetrahedra, wedge (or prismatic) and pyramid elements. This method addressed storage costs on non-tetrahedral elements. However, to accomodate general wedge elements, the method presented in [8] requires the use the rational Low-storage Curvilinear DG (LSC-DG) basis [9], which mitigates the need to explicitly store elemental matrices for wedges in DG methods. This choice of basis for wedge elements in turn requires the use of quadrature-based *a priori* stable variational formulations to maintain energy stability. Additionally, high order convergence under LSC-DG bases requires the use of mesh refinement strategies which result in asymptotically affine elements.

In this paper, we investigate the efficient GPU implementation of nodal DG methods on meshes consisting of both wedge and tetrahedral elements for geophysical applications. The focus is in the efficient solution of wave problems on domains which exhibit either layered features or anisotropy in one direction. Domains of this nature are well-suited to meshes of wedge elements. For example, Lapin et al. [10] utilized meshes consisting of wedges for seismic inversion; however, their focus was in the resolution of interfaces in velocity models, while the focus of this work is the efficient implementation of a time-domain solver for wave propagation and geophysical applications on such meshes. Prismatic meshes have also been adopted in marine applications for resolving topographical variations at ocean floors [11, 12].

In this work, we restrict ourselves to semi-structured prismatic meshes consisting of *vertically mapped* wedges, which are deformed from their reference configuration only along the vertical (extruded) coordinate. By taking advantage of the geometric mappings of such elements, we construct a limited storage, quadrature-free nodal DG method on meshes of tetrahedra and vertically mapped wedge elements. The proposed method addresses issues of robustness to unstructured mesh refinement present for the LSC-DG wedge basis [9, 8], and issues of energy stability associated with mass-lumped approximations.

The paper is structured as follows: we review the continuous and discrete formulations of time domain DG methods for the wave equation in Section 2. In Section 3, we show properties of vertically mapped wedge elements which allow for simplifications of the elemental matrices associated with time domain DG methods, and present a limited-storage method for the acoustic wave equation on such meshes. In Section 5, we present a simple method to construct hybrid meshes of vertically mapped wedges and tetrahedra for layered media. Section 6 gives numerical results confirming the high order convergence of our method, as well as an analysis of the computational performance of the solver at different orders of approximation. Comparisons are drawn between the efficiency of the proposed limited-storage method and LSC-DG on semi-structured prismatic meshes.

2. Discontinuous Galerkin formulation

This work is concerned with numerical solutions of the acoustic wave equation. This equation is posed on a domain Ω with boundary $\partial\Omega$. In first order form, the homogeneous acoustic wave equation is

$$\begin{aligned} \frac{1}{\kappa}\frac{\partial p}{\partial \tau} + \nabla \cdot \boldsymbol{u} &= 0 \\ \rho\frac{\partial \boldsymbol{u}}{\partial \tau} + \nabla p &= 0, \end{aligned}$$

where τ is time, p is pressure, $\boldsymbol{u}$ is the vector velocity (u_x, u_y, u_z), and ρ and κ are density and bulk modulus, respectively. We assume also a triangulation of the domain Ω_h consisting of elements D^k with faces f. Additionally, we assume that ρ and κ are constant over each element, though they may vary spatially from element to element.

Let f be a face between an element D^k and its neighbor D^{k^+}. We define $(p^-, \boldsymbol{u}^-)$ as the restriction of the pressure and velocity respectively on the element D^k to the face f. Similarly, we define test functions

$(\phi^-, \boldsymbol{\psi}^-)$ with support on D^k. Likewise, we define $(p^+, \boldsymbol{u}^+)$ as the restriction of the solution on the element D^{k^+} to f. The jump of p, the normal jump of $\boldsymbol{u}$, and average of $\boldsymbol{u}$ are then defined as

$$[\![p]\!] = p^+ - p^-, \qquad [\![\boldsymbol{u}]\!] = \boldsymbol{u}^+ - \boldsymbol{u}^-, \qquad \{\{\boldsymbol{u}\}\} = \frac{\boldsymbol{u}^+ + \boldsymbol{u}^-}{2}.$$

By defining $\boldsymbol{n}^-$ as the outward unit normal on a given face, the variational formulation for the discontinuous Galerkin method is given over each element D^k as

$$\begin{aligned}
\int_{D^k} \frac{1}{\kappa} \frac{\partial p}{\partial \tau} \phi^- \, \mathrm{d}x &= -\int_{D^k} (\nabla \cdot \boldsymbol{u}) \, \phi^- \, \mathrm{d}x + \int_{\partial D^k} \frac{1}{2} \left(\tau_p [\![p]\!] - \boldsymbol{n} \cdot [\![\boldsymbol{u}]\!] \right) \phi^- \, \mathrm{d}x \\
\int_{D^k} \rho \frac{\partial \boldsymbol{u}}{\partial \tau} \boldsymbol{\psi}^- \, \mathrm{d}x &= -\int_{D^k} \nabla p \cdot \boldsymbol{\psi}^- \, \mathrm{d}x + \int_{\partial D^k} \frac{1}{2} \left(\tau_u [\![\boldsymbol{u}]\!] \cdot \boldsymbol{n}^- - [\![p]\!] \right) \boldsymbol{\psi}^- \cdot \boldsymbol{n}^- \, \mathrm{d}x,
\end{aligned} \tag{1}$$

where $c^2 = \kappa/\rho$ is the speed of sound and $\tau_p = 1/\{\{\rho c\}\}$, $\tau_u = \{\{\rho c\}\}$. This formulation can be shown to be energy stable for any choice of τ_p, τ_u [9], and the specific choice of τ_p, τ_u reduces the numerical flux to the upwind fluxes (as determined by the solution of a Riemann problem) for constant ρ, c. The above formulation is referred to as the "strong" DG formulation [3], where integration by parts is used to ensure all derivatives are applied to the solution variables instead of the weighting functions ϕ and $\boldsymbol{\psi}$.

Finally, for this work, we assume homogeneous Dirichlet boundary conditions $p = 0$ on $\partial\Omega$. These are enforced through reflection conditions at boundary faces $f \in \partial\Omega$

$$p^+\big|_f = -\, p^-\big|_f, \qquad \boldsymbol{u}^+\big|_f = \boldsymbol{u}^-\big|_f.$$

2.1. Discrete formulation

We seek polynomial solutions $(p, \boldsymbol{u})$ over each element D^k such that $(p, \boldsymbol{u})$ satisfy (1). We assume that $p \in V_h\left(D^k\right)$ and $u_x, u_y, u_z \in V_h\left(D^k\right)$, where $V_h\left(D^k\right)$ is a local approximation space on D^k with dimension N_p and basis $\{\phi_i\}_{i=1}^{N_p}$. In Section 3, we make precise the space spanned by this basis, and in Section 3.1 we specify to nodal (Lagrange) polynomial bases.

We assume that each physical element D^k is the image of a reference element $\widehat{D}$, and that a geometric mapping $\boldsymbol{\Phi}^k$ exists such that

$$(x, y, z) = \boldsymbol{\Phi}^k(r, s, t), \qquad (x, y, z) \in D^k, \qquad (r, s, t) \in \widehat{D}.$$

We assume non-degenerate mappings such that the Jacobian J^k of this geometric mapping is greater than zero. Physical approximation spaces are then defined over each element as mappings of a reference approximation space as follows. Let $\{\phi_i\}_{i=1}^{N_p}$ be some basis for a reference approximation space over $\widehat{D}$; then, the physical approximation space over D^k is defined by mappings of this reference space

$$V_h\left(D^k\right) = \operatorname{span}\left\{ \phi_1 \circ \left(\boldsymbol{\Phi}^k\right)^{-1}, \ldots, \phi_{N_p} \circ \left(\boldsymbol{\Phi}^k\right)^{-1} \right\},$$

and each physical basis function is defined as the composition of a reference basis function and the mapping $\boldsymbol{\Phi}^k$ to D^k. Derivatives are computed through reference derivatives and the chain rule.

The discrete formulation of the DG method is given most simply in terms of mass, (weak) derivative, and lift matrices. The mass matrix $\boldsymbol{M}^k$ and face mass matrix $\boldsymbol{M}_f^k$ for the element D^k are defined as

$$\begin{aligned}
\left(\boldsymbol{M}^k\right)_{ij} &= \int_{D^k} \phi_j \phi_i = \int_{\widehat{D}} \phi_j \phi_i J^k \\
\left(\boldsymbol{M}_f^k\right)_{ij} &= \int_{\partial D^{k,f}} \phi_j \phi_i = \int_{\widehat{D}_f} \phi_j \phi_i J_f^k.
\end{aligned}$$

where J_f^k is the Jacobian of the mapping from the face of a reference element $\widehat{D}^k$ to the face of a physical element D_f^k.

We also define weak differentiation matrices $\boldsymbol{S}_x^k, \boldsymbol{S}_y^k, \boldsymbol{S}_z^k$ with entries

$$\left(\boldsymbol{S}_x^k\right)_{ij} = \int_{\widehat{D}} \frac{\partial \phi_j}{\partial x} \phi_i J^k, \qquad \left(\boldsymbol{S}_y^k\right)_{ij} = \int_{\widehat{D}} \frac{\partial \phi_j}{\partial y} \phi_i J^k, \qquad \left(\boldsymbol{S}_z^k\right)_{ij} = \int_{\widehat{D}} \frac{\partial \phi_j}{\partial z} \phi_i J^k.$$

The discrete DG formulation over the kth element may then be given in terms of these matrices

$$\begin{aligned}
\frac{1}{\kappa} \boldsymbol{M}^k \frac{\mathrm{d}\boldsymbol{p}}{\mathrm{d}\tau} &= \left(\boldsymbol{S}_x^k \boldsymbol{U}_x + \boldsymbol{S}_y^k \boldsymbol{U}_y + \boldsymbol{S}_z^k \boldsymbol{U}_z\right) + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{M}_f^k F_p(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+), \\
\rho \boldsymbol{M}^k \frac{\mathrm{d}\boldsymbol{U}_x}{\mathrm{d}\tau} &= \boldsymbol{S}_x^k \boldsymbol{p} + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{n}_x \boldsymbol{M}_f^k F_u(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+) \\
\rho \boldsymbol{M}^k \frac{\mathrm{d}\boldsymbol{U}_y}{\mathrm{d}\tau} &= \boldsymbol{S}_y^k \boldsymbol{p} + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{n}_y \boldsymbol{M}_f^k F_u(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+) \\
\rho \boldsymbol{M}^k \frac{\mathrm{d}\boldsymbol{U}_z}{\mathrm{d}\tau} &= \boldsymbol{S}_z^k \boldsymbol{p} + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{n}_z \boldsymbol{M}_f^k F_u(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+).
\end{aligned}$$

where $\boldsymbol{U}_x, \boldsymbol{U}_y, \boldsymbol{U}_z$ are nodal values for u_x, u_y, u_z. Let D^{k^+} be the neighbor to D^k across a face f; then, $\boldsymbol{p}^-$ are face nodal values for p on D^k, while $\boldsymbol{p}^+$ are face nodal values for p on D^{k^+}. The fluxes F_p, F_u are defined such that

$$\begin{aligned}
\left(\boldsymbol{M}_f^k F_p(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+)\right)_i &= \int_{f_{D^k}} \frac{1}{2} \left(\tau_p \left[\!\left[p\right]\!\right] - \boldsymbol{n}^- \cdot \left[\!\left[\boldsymbol{u}\right]\!\right]\right) \phi_i^- \\
\left(\boldsymbol{n}_i \boldsymbol{M}_f^k F_u(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+)\right)_i &= \int_{f_{D^k}} \frac{1}{2} \left(\tau_u \left[\!\left[\boldsymbol{u}\right]\!\right] \cdot \boldsymbol{n}^- - \left[\!\left[p\right]\!\right]\right) \boldsymbol{\psi}_i^- \boldsymbol{n}_i^-.
\end{aligned}$$

Inverting $\boldsymbol{M}^k$ results in the system of ODEs

$$\begin{aligned}
\frac{1}{\kappa} \frac{\mathrm{d}\boldsymbol{p}}{\mathrm{d}\tau} &= \left(\boldsymbol{D}_x^k + \boldsymbol{D}_y^k + \boldsymbol{D}_z^k\right) \boldsymbol{U}_j + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{L}_f^k F_p(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+), \\
\rho \frac{\mathrm{d}\boldsymbol{U}_x}{\mathrm{d}\tau} &= \boldsymbol{D}_x^k \boldsymbol{p} + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{n}_x \boldsymbol{L}_f^k F_u(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+) \\
\rho \frac{\mathrm{d}\boldsymbol{U}_y}{\mathrm{d}\tau} &= \boldsymbol{D}_y^k \boldsymbol{p} + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{n}_y \boldsymbol{L}_f^k F_u(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+) \\
\rho \frac{\mathrm{d}\boldsymbol{U}_z}{\mathrm{d}\tau} &= \boldsymbol{D}_z^k \boldsymbol{p} + \sum_{f=1}^{N_{\text{faces}}} \boldsymbol{n}_z \boldsymbol{L}_f^k F_u(\boldsymbol{p}^-, \boldsymbol{p}^+, \boldsymbol{U}^-, \boldsymbol{U}^+).
\end{aligned}$$

where we have introduced differentiation matrices $\boldsymbol{D}_x^k, \boldsymbol{D}_y^k, \boldsymbol{D}_z^k$ on the element D^k

$$\boldsymbol{D}_x^k = \left(\boldsymbol{M}^k\right)^{-1} \boldsymbol{S}_x^k, \qquad \boldsymbol{D}_y^k = \left(\boldsymbol{M}^k\right)^{-1} \boldsymbol{S}_y^k, \qquad \boldsymbol{D}_z^k = \left(\boldsymbol{M}^k\right)^{-1} \boldsymbol{S}_z^k$$

and lift matrices $\boldsymbol{L}_f^k$ for each face f of the element D^k

$$\boldsymbol{L}_f^k = \left(\boldsymbol{M}^k\right)^{-1} \boldsymbol{M}_f^k.$$

Using explicit time stepping methods, the solution can be evolved in time through evaluations of local right hand sides for each element, without needing to invert a large global system of equations.

3. Hybrid tet-wedge meshes

In this section, we specialize to meshes whose elements consist of tetrahedra and a restricted class of wedges. We define the space of degree N polynomials over a tetrahedron $\mathcal{T}$ as

$$P^N(\mathcal{T}) = \left\{x^i y^j z^k, \qquad 0 \le i+j+k \le N\right\}, \qquad (x,y,z) \in \mathcal{T}.$$

This space has dimension $N_p = (N+1)(N+2)(N+3)/6$. The space of degree N polynomials for a wedge $\mathcal{W}$ is defined similarly

$$P^N(\mathcal{W}) = \left\{x^i y^j z^k, \qquad 0 \le i+j \le N, \qquad 0 \le k \le N\right\}, \qquad (x,y,z) \in \mathcal{W},$$

with dimension $N_p = (N+1)^2(N+2)/2$. Alternatively, the space of polynomials for the wedge may be defined as the tensor product of degree N polynomials over the reference bi-unit triangle $\widehat{T}$

$$P^N\left(\widehat{T}\right) = \left\{x^i y^j, \quad 0 \le i+j \le N\right\}, \qquad (x,y) \in \widehat{T} = \{-1 \le y \le 1, -1 \le x \le 1-y\}$$

and the space of degree N polynomials $P^N([-1,1])$ over the bi-unit interval.

In this work, we assume low-order geometric mappings for all elements. We refer to such elements as *vertex-mapped*, since specifying physical vertex positions uniquely determines $\mathbf{\Phi}^k$ for $\mathbf{\Phi}^k \in P^1$ for tetrahedral elements and $\mathbf{\Phi}^k \in P^1(\widehat{T}) \times P^1([-1,1])$ for wedge elements. Additionally, we assume high order reference approximation spaces for both wedges and tetrahedra. Under this assumption, the physical approximation space $V_h\left(D^k\right)$ over D^k is

$$V_h\left(D^k\right) = \mathbf{\Phi}^k \circ V_h\left(\widehat{D}\right).$$

which guarantees high order accuracy through reproducibility of polynomials of total degree N [13].

3.1. Nodal bases for tetrahedra and wedges

In this work, we take the basis $\{\phi_i\}_{i=1}^{N_p}$ to be a nodal (Lagrange) polynomial basis $\{\ell_i\}_{i=1}^{N_p}$, such that for a set of points $\{r_i, s_i, t_i\}_{i=1}^{N_p} \in \widehat{D}$,

$$\ell_j(r_i, s_i, t_i) = \delta_{ij}, \qquad 1 \le i, j \le N_p.$$

For quadrature-free DG methods, nodal bases reduce the work involved in the computation of numerical fluxes at face points, since nodal values on a given face can be extracted from the degrees of freedom without additional computation.

For the tetrahedron, nodal bases are defined using Warp and Blend points [14], which are optimized for interpolation quality and numerical stability. For wedges, a nodal basis is defined by taking the tensor product of nodes on the triangular face of the tetrahedron with degree N Gauss-Legendre-Lobatto quadrature nodes in the vertical direction, as shown in Figure 1. Since wedges and tetrahedra are connected through triangular faces, this ensures conformity between wedge and tetrahedral elements. This construction also implies that Lagrange wedge basis functions are tensor products of triangular and one-dimensional Lagrange basis functions $\ell_i^{\text{tri}}(r,s)$ and $\ell^{\text{1D}}(t)$ (respectively)

$$\ell_{ij}(r,s,t) = \ell_i^{\text{tri}}(r,s)\ell_j^{\text{1D}}(t).$$

We index nodal basis functions on the wedge by ij

$$\ell_{ij}(r,s,t), \qquad 0 \le i \le (N+1)(N+2)/2, \qquad 0 \le j \le N,$$

where $(N+1)(N+2)/2$ is the number of nodes in a triangular "slice" of the wedge, i is the index for nodes within a triangular slice, and j enumerates the nodes along the extruded coordinate of the wedge.

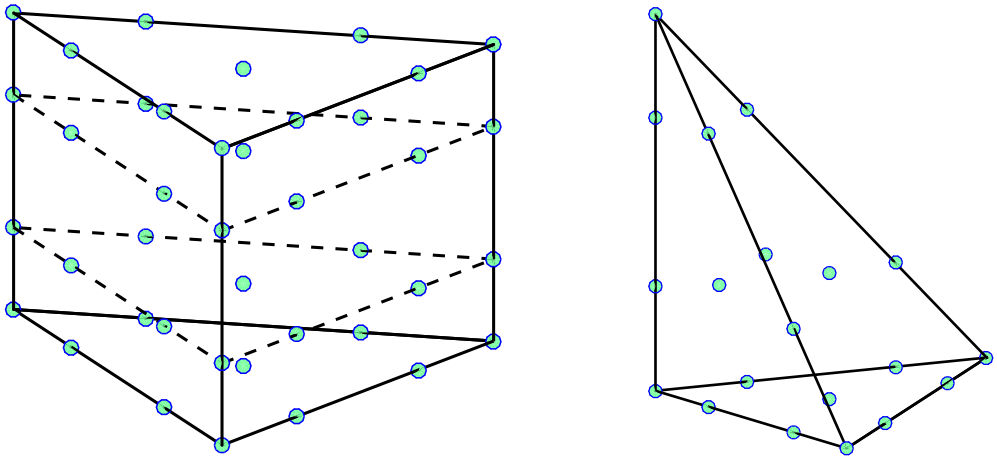

Figure 1: Interpolation points for the wedge and tetrahedron for $N = 3$.

3.2. Geometric factors for vertically mapped wedges

Because any two vertex-mapped tetrahedra are affinely related, geometric factors for tetrahedra are constant from element to element. This results in negligible storage costs, as physical matrices are simply scaled combinations of reference mass, derivative, and lift matrices. However, as noted in [8], polynomial bases which result in diagonal mass matrices for general vertex-mapped wedges are not currently available. As a result, the cost of storing wedge mass matrix inverses or lift matrices can be excessive, especially at high orders of approximation. This is problematic for efficient implementations on many-core architectures (such as GPUs) that typically have very limited memory. In this work, we consider hybrid tet-wedge meshes containing "vertically-mapped" wedges, or wedges which are deformed from their reference configuration along in the vertical (extruded) coordinate. Restricting to such wedges allows us to greatly decrease memory costs for time-domain DG methods.

Physical wedges are images of reference wedges under the local mapping $\boldsymbol{\Phi}^k(r,s,t) = \sum_{i=1}^{6} \boldsymbol{\nu}_i v_i(r,s,t)$, where $\boldsymbol{\nu}_i \in \mathbb{R}^3$ are coordinates of the six vertices on a wedge, and $v_i(r,s,t)$ are low-order vertex basis functions

$$
\begin{aligned}
v_1(r,s,t) &= (1-r-s)(1-t) \\
v_2(r,s,t) &= r(1-t) \\
v_3(r,s,t) &= s(1-t) \\
v_4(r,s,t) &= (1-r-s)t \\
v_5(r,s,t) &= rt \\
v_6(r,s,t) &= st.
\end{aligned}
$$

for reference coordinates $0 \le t \le 1$, $0 \le r, s \le 1$, and $0 \le r+s \le 1$. We define three wedge "vertex pairs" $(\boldsymbol{v}_1, \boldsymbol{v}_4)$, $(\boldsymbol{v}_2, \boldsymbol{v}_5)$, and $(\boldsymbol{v}_3, \boldsymbol{v}_6)$. These correspond to vertices of the reference wedge which have the same r, s coordinates.

The Jacobian of this mapping is defined by

$$
J^k = \det \left[\frac{\partial \boldsymbol{\Phi}^k}{\partial r}, \frac{\partial \boldsymbol{\Phi}^k}{\partial s}, \frac{\partial \boldsymbol{\Phi}^k}{\partial t} \right].
$$

If $\boldsymbol{\nu}_i$ are taken to be the vertex positions of the reference wedge, $\boldsymbol{\Phi}^k = \mathbf{1}$ by the fact that the low order wedge vertex functions $v_i(r,s,t)$ sum to unity.

In this work, we restrict ourselves to *vertically mapped* wedges:

Definition Let D^k be a wedge element; then, D^k is a vertically mapped wedge if the x and y coordinates are identical for each vertex pair $(\boldsymbol{v}_1, \boldsymbol{v}_4)$, $(\boldsymbol{v}_2, \boldsymbol{v}_5)$, and $(\boldsymbol{v}_3, \boldsymbol{v}_6)$.

We note that the definition of vertically mapped wedges allows wedges whose triangular faces are co-planar; however, such elements are not admissible as they have zero volume and thus have $J^k = 0$.

For simplicity, we denote $\frac{\partial r}{\partial x}$ as r_x, and similarly for $\frac{\partial s}{\partial x} = s_x$, $\frac{\partial t}{\partial x} = t_x$, and all other geometric factors. Then, we have the following Lemma:

Lemma 3.1. *For vertically mapped wedges, J^k is constant in the t direction. Additionally, the geometric factors have the following properties*

- *The geometric factors $r_z, s_z = 0$,*
- *The quantities $r_x, r_y, s_x, s_y, \left(t_z J^k\right)$ are constant throughout an element*
- *The quantities $t_x J^k, t_y J^k$ vary only linearly in the t coordinate.*

Proof. We may exploit the fact that the x, y-coordinates of different vertex pairs are identical. For example, the x-coordinates of $\boldsymbol{v}_1$ and $\boldsymbol{v}_4$ are identical, implying that

$$\boldsymbol{v}_1^x v_1(r,s,t) + \boldsymbol{v}_4^x v_4(r,s,t) = \left(\boldsymbol{v}_1^x + \boldsymbol{v}_4^x\right)(1 - r - s).$$

Similarly, the x-coordinates of the $\boldsymbol{v}_2, \boldsymbol{v}_5$ are identical, as are the x-coordinates of $\boldsymbol{v}_3, \boldsymbol{v}_6$, such that

$$\boldsymbol{v}_2^x v_2(r,s,t) + \boldsymbol{v}_5^x v_5(r,s,t) = \left(\boldsymbol{v}_2^x + \boldsymbol{v}_5^x\right) r, \qquad \boldsymbol{v}_3^x v_3(r,s,t) + \boldsymbol{v}_6^x v_6(r,s,t) = \left(\boldsymbol{v}_3^x + \boldsymbol{v}_6^x\right) s.$$

This implies that x is an affine function of r, s, and that $\frac{\partial x}{\partial r}, \frac{\partial x}{\partial s}$ are constant throughout the wedge. As the y-coordinates of each of these vertex pairs is also identical, the same argument implies that $\frac{\partial y}{\partial r}, \frac{\partial y}{\partial s}$ are also constant throughout the wedge. Additionally, since x, y are affine functions of r, s only, this implies $\frac{\partial x}{\partial t}, \frac{\partial y}{\partial t} = 0$.

The z-coordinates of each vertex pair are arbitrary; however, since derivatives of the vertex functions $\frac{\partial v_i}{\partial r}$ and $\frac{\partial v_i}{\partial s}$ are constant in the r, s coordinates, $\frac{\partial z}{\partial r}$ and $\frac{\partial z}{\partial s}$ vary only in the t coordinate. Similarly, since all derivatives $\frac{\partial v_i}{\partial t}$ are constant in the t coordinate, $\frac{\partial z}{\partial t}$ can vary only in the r, s coordinates.

Thus, using that $\frac{\partial x}{\partial t} = \frac{\partial y}{\partial t} = 0$, the expression for J^k simplifies to

$$J^k = \frac{\partial x}{\partial r}\frac{\partial y}{\partial s}\frac{\partial z}{\partial t} - \frac{\partial x}{\partial s}\frac{\partial y}{\partial r}\frac{\partial z}{\partial t}.$$

The partial derivatives $\frac{\partial z}{\partial r}, \frac{\partial z}{\partial s}$ are the only factors which vary in t; since they do not appear in the above expression, J is constant along the t coordinate.

The formulas for geometric factors [2, 3] are also simplified since $x_t, y_t = 0$

$$\begin{aligned}
r_x &= (y_s z_t - z_s y_t)/J^k \qquad & r_y &= -(x_s z_t - z_s x_t)/J^k \qquad & r_z &= (x_s y_t - y_s x_t)/J^k \\
s_x &= -(y_r z_t - z_r y_t)/J^k \qquad & s_y &= (x_r z_t - z_r x_t)/J^k \qquad & s_z &= -(x_r y_t - y_r x_t)/J^k \\
t_x &= (y_r z_s - z_r y_s)/J^k \qquad & t_y &= -(x_r z_s - z_r x_s)/J^k \qquad & t_z &= (x_r y_s - y_r x_s)/J^k.
\end{aligned}$$

Straightforward computations give the remainder of the lemma. □

4. Reducing storage costs for nodal DG methods on semi-structured prismatic meshes

Since wedges are constructed as the tensor product of a triangle and one-dimensional basis, mass lumping techniques are often employed in the extruded direction [2, 15, 16] to reduce storage costs. For a general wedge, mass lumping reduces the cost of explicitly storing $\left(\boldsymbol{M}^k\right)^{-1}$ from $O(N^6)$ to $O(N^5)$, and reduces the cost of explicit storage of the lift matrices $\boldsymbol{L}_f^k$ (under the use of mass-lumped surface quadratures)

from $O(N^5)$ to $O(N^4)$. However, since mass lumping for wedges corresponds to the use of underintegrated quadratures in integrals, they can introduce aliasing errors which result in non-energy stable DG methods. This is discussed in more detail in Section 4.1.

Instead of employing mass lumping, we seek a way to reduce the storage costs of DG for semi-structured prismatic meshes. Given Lemma 3.1, it is possible to reduce the storage of exact quadrature DG methods for meshes with vertically mapped wedges. This method is based on the exact DG formulation (without the use of underintegrated quadratures); however, the implementation of the discrete formulation greatly reduces the storage compared to a straightforward implementation.

4.1. Mass lumping and energy stability

Mass lumped methods take advantage of the fact that Gauss-Legendre-Lobatto nodes are also quadrature nodes, and evaluate entries of the mass matrix using Gauss-Legendre-Lobatto quadrature in the t coordinate. The resulting entries of the wedge mass matrix are

$$\begin{aligned}\left(\boldsymbol{M}^k\right)_{ij,i'j'} &= \int_{\widehat{D}} \ell_{ij}\ell_{i'j'}J = \int_{-1}\int_{\widehat{T}} \ell_{ij}\ell_{i'j'}J(r,s,t) \\ &= \sum_{m=0}^{N}\int_{\widehat{T}} \ell_{ij}(r,s,t_m)\ell_{i'j'}(r,s,t_m)J(r,s,t_m)w_m \\ &= \delta_{jj'}\int_{\widehat{T}} \ell_{ij}(r,s,t_j)\ell_{i'j'}(r,s,t_j)J(r,s,t_j)w_j,\end{aligned}$$

where $\widehat{T}$ is the reference triangle. This mass-lumping procedure results in wedge mass matrices which are block diagonal, where each block is a square matrix of size $\frac{(N+1)(N+2)}{2}$. However, since Gauss-Legendre-Lobatto rules are only exact for polynomials up to degree $(2N-1)$, the use of mass lumping introduces quadrature/aliasing errors. For mass lumped hexes, these errors do not adversely affect energy stability [17]; however, on non-hexahedral elements, these aliasing errors can result in a discrete method which is unstable.

Proofs of energy stability of time-domain DG methods on non-hexahedral elements typically assume exact quadrature, or require *a priori* stable variational formulations in the presence of quadrature errors [9]. Under-integration and mass lumping with the "strong" variational formulation can produce unstable aliasing errors. Recall the local variational formulation in Section 2

$$\begin{aligned}\int_{D^k}\frac{1}{\kappa}\frac{\partial p}{\partial \tau}\phi^-\,\mathrm{d}x &= -\int_{D^k}\nabla\cdot\boldsymbol{u}\phi^-\,\mathrm{d}x + \int_{\partial D^k}\frac{1}{2}\left(\tau_p\left[\!\left[p\right]\!\right] - \boldsymbol{n}\cdot\left[\!\left[\boldsymbol{u}\right]\!\right]\right)\phi^-\,\mathrm{d}x \\ \int_{D^k}\rho\frac{\partial \boldsymbol{u}}{\partial \tau}\boldsymbol{\psi}^-\,\mathrm{d}x &= -\int_{D^k}\nabla p\cdot\boldsymbol{\psi}^-\,\mathrm{d}x + \int_{\partial D^k}\frac{1}{2}\left(\tau_u\left[\!\left[\boldsymbol{u}\right]\!\right]\cdot\boldsymbol{n}^- - \left[\!\left[p\right]\!\right]\right)\boldsymbol{\psi}^-\cdot\boldsymbol{n}^-\,\mathrm{d}x.\end{aligned}$$

If integration by parts holds at the discrete level, the above formulation is discretely equivalent to the "skew-symmetric" formulation

$$\begin{aligned}\int_{D^k}\frac{1}{\kappa}\frac{\partial p}{\partial \tau}\phi^-\,\mathrm{d}x &= \int_{D^k}\boldsymbol{u}\cdot\nabla\phi^-\,\mathrm{d}x + \int_{\partial D^k}\frac{1}{2}\left(\tau_p\left[\!\left[p\right]\!\right] - 2\boldsymbol{n}\cdot\{\{\boldsymbol{u}\}\}\right)\phi^-\,\mathrm{d}x \\ \int_{D^k}\rho\frac{\partial \boldsymbol{u}}{\partial \tau}\boldsymbol{\psi}^-\,\mathrm{d}x &= -\int_{D^k}\nabla p\cdot\boldsymbol{\psi}^-\,\mathrm{d}x + \int_{\partial D^k}\frac{1}{2}\left(\tau_u\left[\!\left[\boldsymbol{u}\right]\!\right]\cdot\boldsymbol{n}^- - \left[\!\left[p\right]\!\right]\right)\boldsymbol{\psi}^-\cdot\boldsymbol{n}^-\,\mathrm{d}x,\end{aligned}$$

then energy stability follows by taking $\phi^- = p$ and $\boldsymbol{\psi}^- = \boldsymbol{u}$ (see, for example [3, 18]). Pulling the time derivative out of the left hand side term and summing both equations over all elements D^k yields

$$\begin{aligned}&\frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}\tau}\sum_{D^k\in\Omega_h}\left\{\int_{D^k}\left(\frac{1}{\kappa}p^2 + \rho\boldsymbol{u}^2\right)\right\} = \\ &\frac{1}{2}\sum_{D^k\in\Omega_h}\int_{\partial D^k}\tau_p\left[\!\left[p\right]\!\right]p^- + \tau_u\left[\!\left[\boldsymbol{u}\right]\!\right]\cdot\boldsymbol{n}^-\boldsymbol{u}^-\cdot\boldsymbol{n}^- - 2\boldsymbol{n}\cdot\{\{\boldsymbol{u}\}\}p^- - \left[\!\left[p\right]\!\right]\boldsymbol{u}^-\cdot\boldsymbol{n}^-\,\mathrm{d}x. \qquad (2)\end{aligned}$$

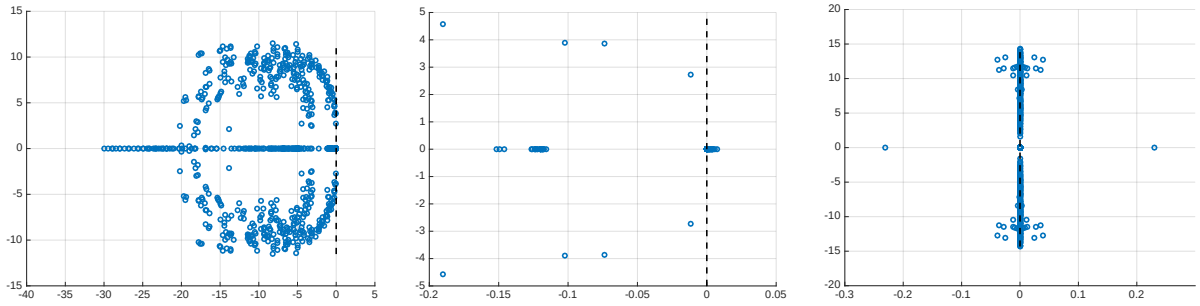

(a) Upwind flux (b) Upwind flux (zoom) (c) Central flux

Figure 2: Eigenvalues for a mass-lumped DG discretization on wedges with $N = 2$. Unlike non-mass lumped discretizations, there exist eigenvalues with positive real parts, implying that energy stability is lost due to errors introduced by underintegration.

The surface terms may be rewritten as sums over unique faces $f \in \Gamma_h$, where Γ_h denotes the set of unique faces in Ω_h. Expanding out the face integrands in terms of $\boldsymbol{n}^-$, $(p^-, \boldsymbol{u}^-)$, and $(p^+, \boldsymbol{u}^+)$ shows that the terms which are independent of τ_p, τ_u cancel. The remaining penalization terms then reduce to $-\tau_p \left[\!\left[p\right]\!\right]^2$ and $-\tau_u \left[\!\left[\boldsymbol{u}_n\right]\!\right]^2$ over each face, and inserting them into (2) yields

$$\frac{\mathrm{d}}{\mathrm{d}\tau} \sum_{D^k \in \Omega_h} \left(\int_{D^k} \left(\frac{1}{\kappa} p^2 + \rho \boldsymbol{u}^2 \right) \right) = - \sum_{f \in \Gamma_h} \int_{\partial f} \left(\tau_p \left[\!\left[p\right]\!\right]^2 + \tau_u \left[\!\left[\boldsymbol{u}_n\right]\!\right]^2 \right) \leq 0.$$

This implies that for $\tau_p, \tau_u \geq 0$, DG methods exhibit dissipative behavior, where unresolvable components of the solution corresponding to high frequency modes are damped [19, 20]. More precisely, after discretizing in space using DG, the PDE reduces to a system of ODEs for the nodal values $\boldsymbol{U}$

$$\frac{\mathrm{d}\boldsymbol{U}}{\mathrm{d}\tau} = \boldsymbol{A}\boldsymbol{U},$$

where we refer to $\boldsymbol{A}$ as the DG right-hand side matrix. For $\tau_p, \tau_u > 0$, we refer to the flux as an *upwind flux*. For upwind fluxes, the eigenvalues of $\boldsymbol{A}$ contain negative real parts, indicating dissipation in time of the solution. Additionally, for $\tau_p = \tau_u = 0$, the DG formulation reduces to that of a central flux, resulting in a non-dissipative method where the eigenvalues of $\boldsymbol{A}$ are purely imaginary. In the following experiments, we build $\boldsymbol{A}$ for $\rho = c^2 = 1$, homogeneous Dirichlet boundary conditions, and mesh of 16 wedges whose vertices are randomly perturbed in the vertical direction to ensure non-constant J^k.

For Discontinuous Galerkin spectral element methods (DG-SEM), which utilize mass lumping on quadrilateral or hexahedral elements, integration by parts holds under Gauss-Legendre-Lobatto quadrature and the strong and skew-symmetric formulation are discretely equivalent [17, 21]. This equivalence relies on the diagonal nature of the mass matrix and the exactness of Gauss-Legendre-Lobatto quadrature for degree $(2N - 1)$ polynomials. Unfortunately, the analogous statement does not hold true for mass lumped wedges. Figure 2 shows the spectra of the DG right hand side operator resulting from a discretization of the strong formulation using mass lumping on a mesh with vertically mapped wedges. For both upwind and central fluxes, eigenvalues with real positive part are present, implying that the formulation is not energy stable. In contrast, employing a non-mass lumped discretization is equivalent to using an exact quadrature rule, such that continuous integration by parts implies discrete equivalence of the strong formulation and energy-stable skew-symmetric formulation. This is reflected in Figure 3, which shows the eigenvalues of the DG operator under our method with exact quadrature.

4.2. Mass, differentiation, and lift matrices

Since mass lumping on wedges produces DG methods which are not energy stable, we seek to reduce the storage costs for exact quadrature DG methods on wedge meshes, taking advantage of the fact that for

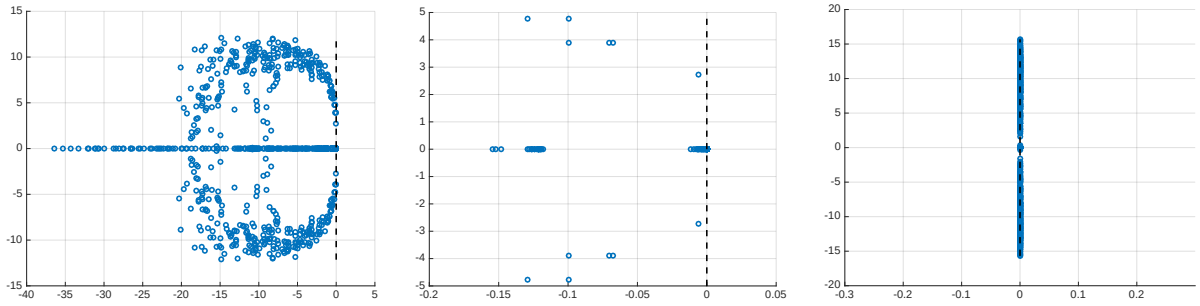

(a) Upwind flux (b) Upwind flux (zoom) (c) Central flux

Figure 3: Eigenvalues for the exact quadratureDG discretization on wedges with $N = 2$. All eigenvalues have non-positive real part (up to machine precision), implying that the formulation is energy stable.

vertically mapped wedges, J^k is constant in t. Since wedge nodal basis functions can be defined as the tensor product of triangular and one-dimensional nodal basis functions, the mass matrix over an element D^k has entries

$$\left(\boldsymbol{M}^k\right)_{ij,i'j'} = \int_{\widehat{D}} \ell_{ij}\ell_{i'j'}J = \int_{-1}^{1} \ell_j^{\mathrm{1D}}(t)\ell_{j'}^{\mathrm{1D}}(t) \int_{\widehat{T}} \ell_i^{\mathrm{tri}}\ell_{i'}^{\mathrm{tri}}J(r,s).$$

This implies that $\boldsymbol{M}^k$ is expressible as

$$\boldsymbol{M}^k = \boldsymbol{M}^{\mathrm{tri},k} \otimes \widehat{\boldsymbol{M}}^{\mathrm{1D}}$$

where the entries of the triangle mass matrix $\boldsymbol{M}^{\mathrm{tri},k}$ for a given element are

$$\left(\boldsymbol{M}^{\mathrm{tri},k}\right)_{ij} = \int_{\widehat{T}} \ell_i^{\mathrm{tri}}(r,s)\ell_j^{\mathrm{tri}}(r,s)J(r,s).$$

We include superscripts k in $\boldsymbol{M}^{\mathrm{tri},k}$ to emphasize that the triangle mass matrices are distinct from element to element.

Face mass matrices under vertically mapped wedges also possess a Kronecker product structure. The entries of face mass matrices for a face f of an element D^k are given by

$$\left(\boldsymbol{M}_f^k\right)_{ij,i'j'} = \int_f \ell_{ij}\ell_{i'j'}J_f^k,$$

where J_f^k is the Jacobian of the mapping from a reference triangle or quadrilateral to the physical triangular/quadrilateral face of an element. For triangular faces, J_f^k is constant, since any two planar triangles are affinely related. Since wedge nodes include nodes on each of the triangular faces, entries of the triangular face mass matrix $\left(\boldsymbol{M}_f^k\right)_{ij,i'j'}$ are nonzero only for nodes where $j, j' = 0$ or $j, j' = N$.

$$\boldsymbol{M}_f^k = J_f^k \widehat{\boldsymbol{M}}^{\mathrm{tri}} \otimes \boldsymbol{e}$$

where $\widehat{\boldsymbol{M}}^{\mathrm{tri}}$ is the reference triangle mass matrix, and $\boldsymbol{e} \in \mathbb{R}^{N+1}$ is the first canonical basis vector for the triangle face corresponding to $t = -1$ and the last canonical basis vector for the triangular face corresponding to $t = 1$, respectively.

For general quadrilateral faces, J_f^k is not constant. However, for vertically mapped wedges, quadrilateral vertices are displaced only in the vertical direction, and quadrilateral faces are planar with constant normals.

Using similar arguments as in Lemma 3.1, it is possible to show that J_f^k is constant in the extruded (vertical) t coordinate, and varies linearly in the horizontal coordinate. The mass matrix for a quadrilateral face also possesses a Kronecker product structure

$$\boldsymbol{M}_f^k = \tilde{\boldsymbol{M}}_{\text{edge}}^{\text{tri},k} \otimes \widehat{\boldsymbol{M}}^{\text{1D}}$$

where $\tilde{\boldsymbol{M}}_{\text{edge}}^{\text{tri},k}$ is the face (edge) mass matrix for a triangle, which is computed using the restriction of the quadrilateral face Jacobian J_f^k to the edge

$$\left(\tilde{\boldsymbol{M}}_{\text{edge}}^{\text{tri},k}\right)_{ij} = \int_{-1}^{1} \ell_i^{\text{tri}} \ell_j^{\text{tri}} J_f^k.$$

Since J_f^k varies along the edge, $\tilde{\boldsymbol{M}}_{\text{edge}}^{\text{tri},k}$ is distinct from face to face.

The tensor-product nature of the face mass matrices simplifies the form of the lift matrix. For a triangular face f, the lift matrix $\boldsymbol{L}_f^k$ reduces to

$$\boldsymbol{L}_f^k = \left(\boldsymbol{M}^k\right)^{-1} \boldsymbol{M}_f^k = J_f^k \left(\boldsymbol{M}^{\text{tri},k}\right)^{-1} \widehat{\boldsymbol{M}}^{\text{tri}} \otimes \left(\widehat{\boldsymbol{M}}^{\text{1D}}\right)^{-1} \boldsymbol{e},$$

while the lift matrix for a quadrilateral face f is

$$\boldsymbol{L}_f^k = \left(\boldsymbol{M}^{\text{tri},k}\right)^{-1} \tilde{\boldsymbol{M}}_{\text{edge}}^{\text{tri},k} \otimes \boldsymbol{I}.$$

This implies that, for vertically mapped wedges, lift matrices for quadrilateral faces are always block diagonal.

We refer to the matrix $\left(\boldsymbol{M}^{\text{tri},k}\right)^{-1} \widehat{\boldsymbol{M}}^{\text{tri}}$ as the triangular lift $\boldsymbol{L}^{\text{tri},k}$. Both lift matrices for triangular and quadrilateral faces are tensor products of the triangular lift matrix and a one-dimensional matrix. The triangular lift matrix also appears again in the expression for differentiation matrices.

Since nodal basis functions for the wedge are constructed in a tensor-product fashion, reference differentiation matrices possess a Kronecker product structure. We define $\boldsymbol{D}_r^{\text{tri}}, \boldsymbol{D}_s^{\text{tri}}$ to be differentiation matrices for nodal basis functions defined on the triangle (using Warp and Blend nodes for the face of a tetrahedron). Likewise, we define $\boldsymbol{D}_t^{\text{1D}}$ as the differentiation matrix for a Lagrange basis at degree N Gauss-Legendre-Lobatto nodes in one dimension. Then, the reference wedge differentiation matrices are

$$\boldsymbol{D}_r = \boldsymbol{D}_r^{\text{tri}} \otimes \boldsymbol{I}^{\text{1D}}, \qquad \boldsymbol{D}_s = \boldsymbol{D}_s^{\text{tri}} \otimes \boldsymbol{I}^{\text{1D}}, \qquad \boldsymbol{D}_t = \boldsymbol{I}^{\text{tri}} \otimes \boldsymbol{D}_t^{\text{1D}},$$

where $\boldsymbol{I}^{\text{tri}}, \boldsymbol{I}^{\text{1D}}$ are square identity matrices of size $(N+1)(N+2)/2$ and $(N+1)$, respectively.

Weak differentiation matrices also exhibit a Kronecker product structure for vertically mapped wedges. Lemma 3.1 gives that r_x, s_x, r_y, s_y are constant within an element, such that

$$\left(\boldsymbol{S}_x^k\right)_{ij,i'j'} = \int_{\widehat{D}} \frac{\partial \ell_{i'j'}}{\partial x} \ell_{ij} J^k = r_x \int_{\widehat{D}} \frac{\partial \ell_{i'j'}}{\partial r} \ell_{ij} + s_x \int_{\widehat{D}} \frac{\partial \ell_{i'j'}}{\partial s} \phi_{ij} J^k + \int_{\widehat{D}} \frac{\partial \ell_{i'j'}}{\partial t} \ell_{ij} t_x J^k.$$

Noting from Lemma 3.1 that $t_x J^k$ is a polynomial of degree 1 in t and that $\frac{\partial \ell_{i'j'}}{\partial t}$ is a polynomial of degree $(N-1)$, their product is a polynomial of degree N. Since multiplication by the mass is equivalent to integration of a degree N polynomial against elements of the basis ℓ_{ij}, we have that

$$\int_{\widehat{D}} t_x J^k \frac{\partial \ell_{i'j'}}{\partial t} \ell_{ij} = \left(\boldsymbol{M}^{\text{tri},k} \otimes \widehat{\boldsymbol{M}}^{\text{1D}} \text{diag}(t_x J^k) \boldsymbol{D}_t^{\text{1D}}\right)_{ij,i'j'},$$

where $\text{diag}(t_x J^k)$ is a diagonal matrix of $t_x J^k$ at Gauss-Legendre-Lobatto nodes in the t coordinate. The same argument can be made for $\boldsymbol{S}_y^k$, implying that

$$\begin{aligned}
\boldsymbol{D}_x^k &= \left(r_x \boldsymbol{D}_r^{\text{tri}} + s_x \boldsymbol{D}_s^{\text{tri}}\right) \otimes \boldsymbol{I}^{\text{1D}} + \boldsymbol{L}^{\text{tri},k} \otimes \text{diag}(t_x J^k) \boldsymbol{D}_t^{\text{1D}} \\
\boldsymbol{D}_y^k &= \left(r_y \boldsymbol{D}_r^{\text{tri}} + s_y \boldsymbol{D}_s^{\text{tri}}\right) \otimes \boldsymbol{I}^{\text{1D}} + \boldsymbol{L}^{\text{tri},k} \otimes \text{diag}(t_y J^k) \boldsymbol{D}_t^{\text{1D}}.
\end{aligned}$$

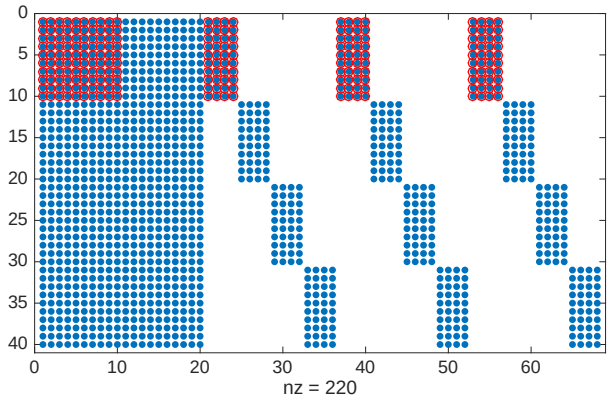

Figure 4: Spy plot of a lift matrix for a vertically mapped wedge. Explicitly stored entries are circled in red. The axes correspond to row and column indices.

The differentiation matrix $\boldsymbol{D}_z^k$ also possesses a tensor product structure; additionally, since $r_z, s_z = 0$ and t_z is constant for vertically mapped wedges, the expression for $\boldsymbol{D}_z^k$ simplifies to

$$\boldsymbol{D}_z^k = \left(\boldsymbol{M}^{\text{tri},k}\right)^{-1} \widehat{\boldsymbol{M}}^{\text{tri}} \otimes t_z J^k \boldsymbol{D}_t^{\text{1D}} = \boldsymbol{L}^{\text{tri},k} \otimes t_z J^k \boldsymbol{D}_t^{\text{1D}}.$$

If the mapping is affine, the triangular lift matrix $\boldsymbol{L}^{\text{tri},k}$ reduces to $J^k \boldsymbol{I}^{\text{tri}}$, weak differentiation matrices reduce to scalings of nodal differentiation matrices [22].

Finally, since the method in this paper uses exact quadrature, all the terms in the DG formulation (1) are computed exactly and theoretical results for standard DG methods are automatically inherited, such as high order accuracy and energy stability [23, 3].

4.3. Reduced storage costs

For general vertex-mapped hexahedra, pyramids, and tetrahedra, there exist bases or stable mass-lumped approximations which yield low-storage mass matrices [2, 24]. However, unlike other element types, time domain DG methods on wedges (using polynomial bases) require explicit storage of matrices for each element. Naive storage of elemental matrices can increase the available memory required from $O(N^3)$ to $O(N^6)$ for each element in three dimensions. These large storage costs can greatly limit the admissible problem sizes on GPUs and modern accelerator architectures, which tend to possess limited on-device memory. There does not appear to be a polynomial basis which yields (for non-affine wedges) lift matrices which requires $O(N^3)$ or less storage per element [8]. However, by assuming vertically mapped wedges, these storage costs may be reduced to a more acceptable level.

Restricting to vertically mapped wedge elements exposes a Kronecker product structure in both differentiation and lift matrices involving the triangular lift matrix $\boldsymbol{L}^{\text{tri},k} = \left(\boldsymbol{M}^{\text{tri},k}\right)^{-1} \widehat{\boldsymbol{M}}^{\text{tri}}$. Since $\left(\boldsymbol{M}^{\text{tri},k}\right)^{-1}$ varies depending on the local geometric mapping, we explicitly store $\boldsymbol{L}^{\text{tri},k}$ in order to avoid solving a matrix equation. Storage of $\boldsymbol{L}^{\text{tri},k}$ requires $O(N^4)$ storage per element. The lift matrix for quadrilateral faces of the wedge also requires the matrices $\left(\boldsymbol{M}^{\text{tri},k}\right)^{-1} \widetilde{\boldsymbol{M}}_{\text{edge}}^{\text{tri},k}$, for which explicit storage requires only $O(N^3)$ storage per element. Thus, while we have increased the asymptotic storage cost to $O(N^4)$, we only store one additional matrix of this size per element.

Figure 4 shows the storage required for the lift matrix as compared to the storage of the full lift matrix. The matrix shown is a column-wise concatenation of lift matrices for both triangular and quadrilateral faces. The first $(N+1)(N+2)/2$ columns correspond to the lift matrix for the bottom triangular face, while the next $(N+1)(N+2)/2$ columns correspond to the lift matrix for the top triangular face. The remaining three blocks of $(N+1)^2$ columns correspond to lift matrices for each of the three quadrilateral faces of the wedge. Interestingly, one can show that, for vertically mapped wedges, lift matrices under mass lumping have the same storage requirements.

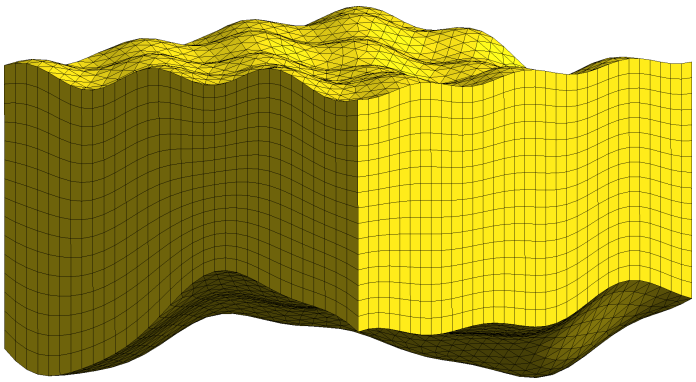

Figure 5: Example of meshing a layer defined by two arbitrary surfaces using wedges.

5. Hybrid tet-wedge meshes for complex interfaces

We are interested in resolving geometries which contain complex interfaces with separate distinct layers, such as those found in geophysical applications (ocean-earth interfaces and layers of strata). The approach taken in this work is to develop meshes consisting of wedge elements which conform to interfaces. These approaches are also applicable to high order models of atmospheric or coastal flow, which often represent earth-air or ocean-air interfaces based on unstructured meshes of wedges [25, 26]. Other techniques for representing complex interfaces include coordinate transformations, such as the sigma transform [27]. Typical applications of the sigma coordinate method map a domain with smooth boundaries to a Cartesian domain, where the solution can be determined by solving a modified PDE using numerical methods for structured grids. In contrast, directly approximating the domain using semi-structured prismatic meshes can handle more general geometries, including non-smooth interfaces. Coupling to tetrahedral meshes is also made simpler through the use of vertically mapped wedge meshes rather than sigma transformations.

Constructing multi-layered wedge meshes can be done in layer-by-layer fashion. We assume that a single layer domain Ω_L is defined as

$$\Omega_L = \{(x, y) \in [a, b] \times [c, d], \quad z_{\text{bottom}}(x, y) \leq z \leq z_{\text{top}}(x, y)\},$$

where $z_{\text{top}}(x, y), z_{\text{bottom}}(x, y)$ are surfaces defining the top and bottom boundaries of the layer. A surface triangulation is then constructed which resolves both top and bottom layers. The final step is to produce a quasi-uniform mesh which resolves the interior of the layer. To do so, we construct a single-layer reference wedge mesh by extruding the surface triangulation and interpolate the surfaces $z_{\text{top}}(x, y)$ and $z_{\text{top}}(x, y)$ with the vertices at the top and bottom of the reference wedge mesh, respectively. This produces a linear interpolant between the top and bottom surfaces, and evaluating this interpolant at a point z and at the (x, y) coordinates of the surface triangulation vertices produces the coordinates of an intermediate layer. Evaluating this interpolant at equispaced points in the vertical coordinate of the reference wedge mesh produces multiple layers, which may then be connected to produce a quasi-uniform wedge mesh which conforms to $z_{\text{top}}(x, y)$ and $z_{\text{bottom}}(x, y)$. An example of such a wedge mesh is shown in Figure 5.

The construction of meshes for layered media can also be done using extrusions of quadrilaterals, resulting in more efficient hexahedral grids. However, an advantage of wedge meshes is their ability to conform to surfaces of tetrahedral meshes. This is useful in some geophysical applications, where meshes which conform to interfaces in high contrast media are desirable for capturing sharp transitions between areas of differing wavespeed. In these cases, it is not typically possible to partition strata of the earth into distinct layers, requiring instead more flexible tetrahedral mesh generation algorithms to produce interface-conforming meshes. However, wedge meshes may be still be used to capture earth-ocean interfaces or layered structures if present, while tetrahedral meshes are used to capture non-layered high contrast interfaces within the earth. Figure 6 shows an example of such a mesh, constructed using the above procedure for layered wedge meshes and the Tetgen mesh generation package [28].

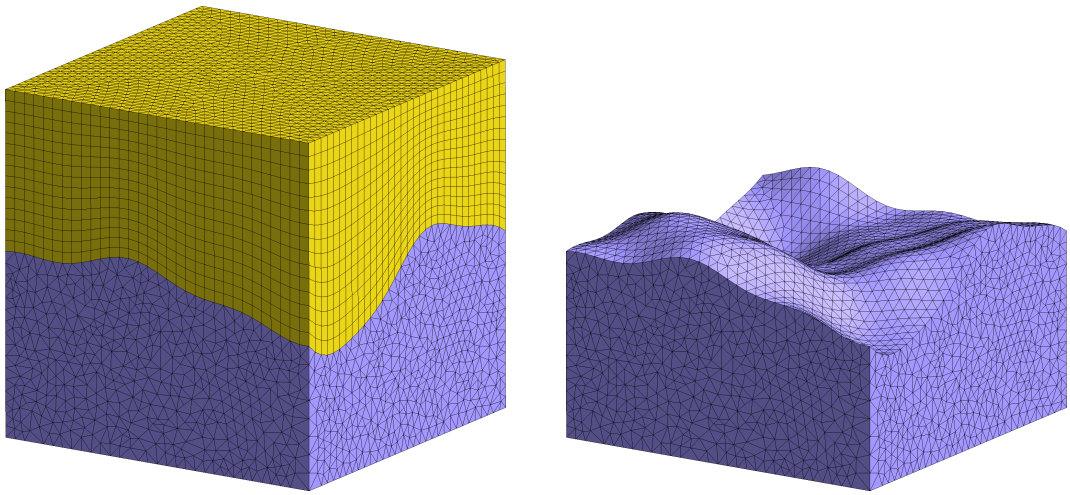

Figure 6: Hybrid wedge-tet meshes for capturing interfaces.

6. Numerical experiments

In the following sections, we present numerical experiments which confirm the high order accuracy and efficiency of the presented methods on many-core architectures. We also present experiments which highlight differences between between different approaches to high order DG methods on hybrid meshes.

A GPU-accelerated high order DG method on hybrid meshes was presented in [8]; however the method did not inherit several key features of high order nodal DG methods on tetrahedral meshes. First, the use of mass lumping and Gauss-Legendre-Lobatto underintegration was shown to decrease convergence rates by half an order in comparison to the use of full quadrature. Secondly, due to the use of a rational Low-Storage Curvilinear DG (LSC-DG) basis [29, 9], high order convergence rates were observed but were contingent on refinement strategies resulting in asymptotically affine elements. The use of LSC-DG reduces storage costs on general vertex-mapped wedges, but results in a method which may not be robust to mesh perturbation, implying that high order rates of convergence could be lost for non-nested or unstructured mesh refinement. Finally, the use of LSC-DG necessitated the explicit use of quadrature on faces for energy stability, and requires the evaluation of solutions at quadrature points on triangular and quadrilateral faces. This procedure introduces additional computational work with $O(N^5)$ complexity compared to quadrature-free nodal DG methods.

The method presented addresses each of these issues for vertically mapped wedges, resulting in a quadrature-free nodal DG method which is both high order accurate method and robust to vertical mesh perturbation. Discrete timestep restrictions were determined using geometric factors and constants in trace and Markov inequalities [8]. Experiments were performed on an Nvidia GTX 980 GPU with a cross-platform implementation leveraging the OCCA framework [30].

6.1. Convergence tests

A significant advantage of high order approximations is that, compared to a low order approximation, far fewer degrees of freedom are required to approximate a sufficiently regular function to a given tolerance. This is reflected in the convergence rates for high order methods: for a degree N order polynomial approximation, the L^2 error converges with rate $O\left(h^{N+1}\right)$, where h is the mesh size. DG methods with upwind numerical fluxes converge at this optimal rate on a large class of meshes [31], though in general one can expect only a $O\left(h^{N+1/2}\right)$ rate of convergence [32].

It is also possible to diminish the convergence rate of DG methods through underintegration. For example, mass-lumped DG-SEM methods are observed to converge in the L^2 norm at rate $O\left(h^{N+1/2}\right)$ [33, 8]. This

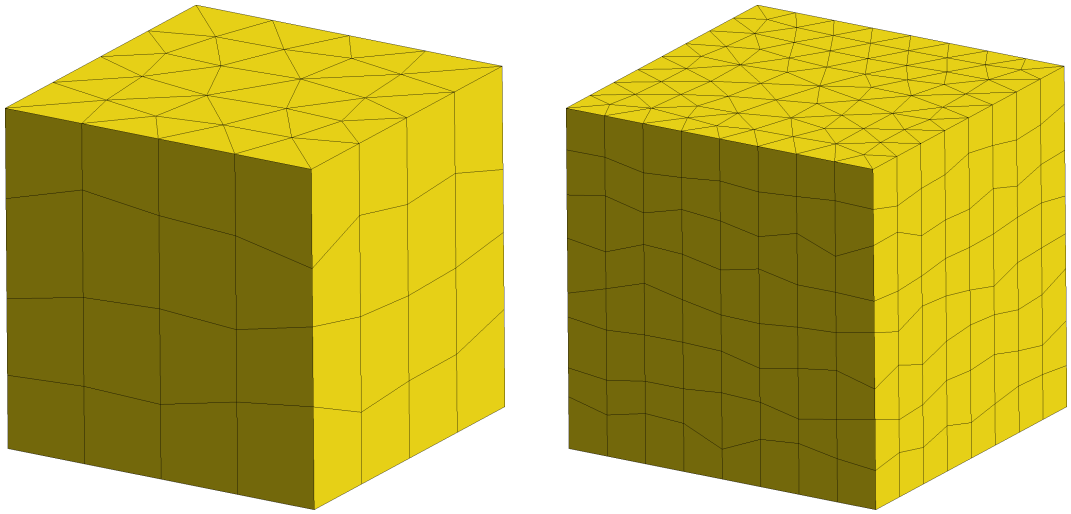

Figure 7: Perturbed vertically mapped wedge meshes used in unstructured mesh refinement tests.

is observed to improve to $O\left(h^{N+1}\right)$ if exact quadrature rules are used, at the cost of introducing an extra interpolation step to compute solution values on faces.

In contrast, the reduced storage DG method in this work avoids both underintegration and extra interpolation steps. Figure 8 shows the convergence of the L^2 error for an exact pressure solution

$$p(x, y, z, t) = \cos\left(\frac{\pi x}{2}\right) \cos\left(\frac{\pi y}{2}\right) \cos\left(\frac{\pi z}{2}\right) \cos\left(\frac{\sqrt{3}\pi t}{2}\right)$$

under mesh refinement, and optimal rates of convergence (Table 1) are observed.

To verify the robustness of our method with respect to unstructured mesh refinement, we examine convergence also on a sequence of vertically mapped perturbed wedge meshes. Each wedge mesh is generated by extruding an unstructured surface triangulation vertically into several layers, then perturbing the vertical coordinates of each wedge randomly (as shown in Figure 7). Optimal rates of convergence of L^2 error are observed (Figure 8 and Table 1) for both structured and unstructured mesh refinement.

We also compare the L^2 error of our limited storage implementation of standard DG to that of LSC-DG on each set of meshes in Figure 8. On affine meshes, LSC-DG reduces to a standard DG method, and we observe optimal convergence for both methods. The error for LSC-DG is slightly less than that for DG at lower orders, due to the fact that the initial condition is computed as an L^2 projection instead of through Lagrange interpolation (though by $N = 3$, this difference is negligible). For perturbed meshes, L^2 errors for LSC-DG begin to stall at small mesh sizes h, due to the fact that under perturbed mesh refinement, elements are not asymptotically affine.

To emphasize that the convergence of LSC-DG can stagnate under specific types of mesh refinement, we consider a sequence of vertically mapped wedge meshes based on the quadrilateral meshes used in [34, 9]. Each mesh in this sequence is constructed through a self-similar refinement pattern. Figure 9 shows an example of one such mesh, as well as the L^2 errors of DG and LSC-DG on such meshes. While L^2 errors for DG are similar to those observed under structured mesh refinement, errors for LSC-DG stagnate under mesh refinement.

6.2. Computational implementation

Time-domain DG solvers using explicit time stepping typically advance the solution in time in two stages - a step which computes the evaluation of the right hand side, and one which computes the solution at the

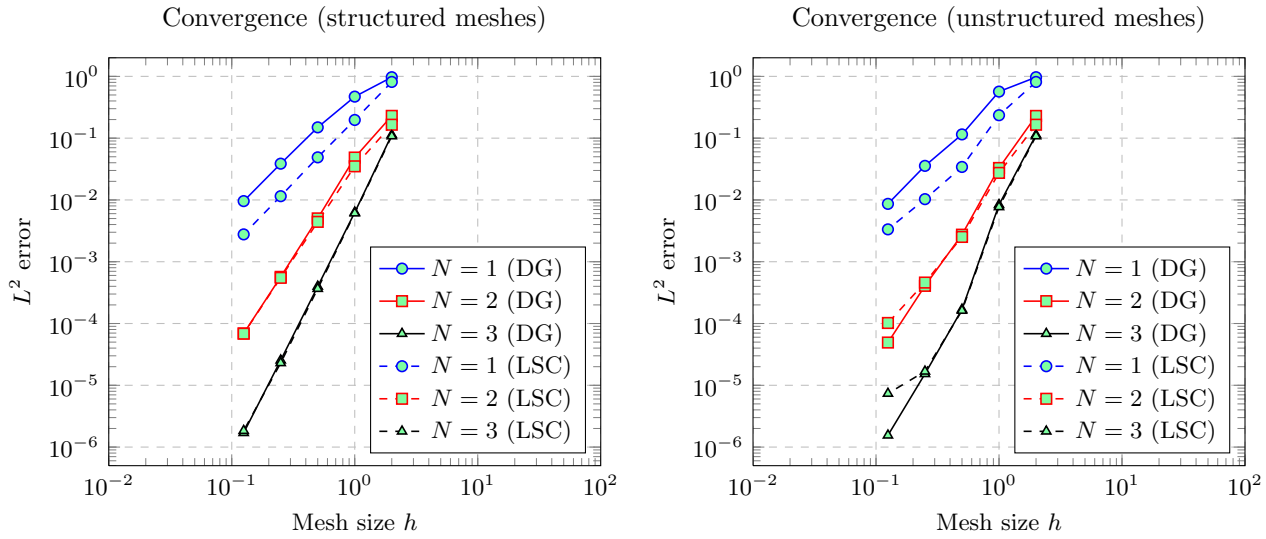

Figure 8: Convergence of the L^2 error for LSC-DG and DG under structured and unstructured wedge mesh refinement.

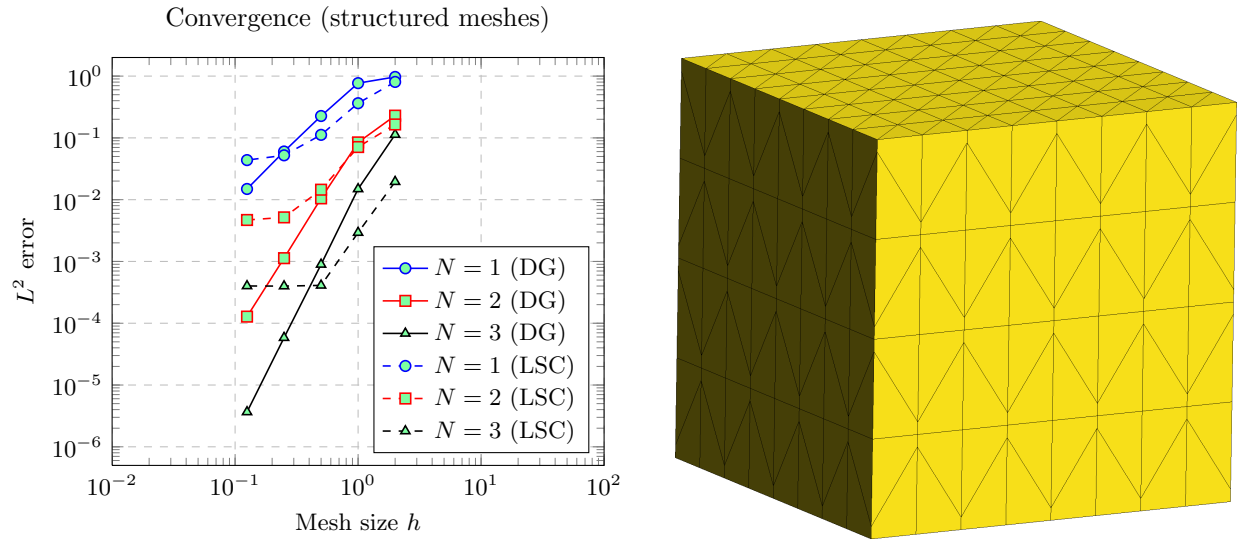

Figure 9: L^2 errors under mesh refinement for Arnold-type perturbed meshes.

	$N = 1$	$N = 2$	$N = 3$
Structured	2.01	3.15	3.97
Unstructured	1.72	2.9	4.42
Arnold-type	1.9	3.13	3.99

Table 1: Observed rates of convergence of L^2 error for DG under various types of mesh refinement.

next timestep using a time integrator. Following Klöckner et al. [4], we implement the above steps on a GPU by dividing this work into three distinct kernels for each type of element (wedge, tetrahedron):

- A volume kernel, which applies differentiation matrices to compute contributions to the right hand side resulting from volumetric integrals in the variational formulation.
- A surface kernel, which computes both numerical fluxes and applies the lift matrix to compute contributions to the right hand side resulting from surface integrals in the variational formulation.
- An update kernel, which updates the solution in time.

For the update kernel, we use a 3rd order multi-rate Adams Bashforth time integrator in this work [35], though any standard explicit time marching scheme may be used. The implementation of wedge update kernels and tetrahedral volume, surface, and update kernels follows [6]. Since the computational cost of the update kernel is minimal for quadrature-free nodal DG methods, we do not analyze its computational performance in this work. Instead, we focus on the implementation of volume and surface kernels for the vertically mapped wedge elements.

Finally, achieving good computational performance on GPUs typically requires tuning computational parameters to ensure work is distributed as evenly as possible among available GPU resources. The work is divided such that a thread block processes computation for one or more elements. We assign computational parameters K_V, K_S, and K_U, which we refer to respectively as the number of elements processed by the volume, surface, and update kernels in each thread block. These parameters can be tuned such that the number of total threads processed is made close to a multiple of 32, the number of threads which run concurrently within one thread block. For all following computational results, K_V, K_S, and K_U have been optimized to minimize runtime.

6.3. GPU implementation

In this section, we discuss parallelization strategies used in the implementation of volume and surface kernels on GPUs. For the volume surface kernel, we parallelize over nodes in a triangular slice of the wedge, marching through slices of the wedge in serial. This strategy takes advantage of the triangle-line tensor product structure of local matrices by decomposing computation into multiplication by one-dimensional and triangular matrices. One-dimensional matrices are applied for each thread (corresponding to a node in a given triangular slice) by loading entries of the matrix from memory and computing the matrix product with data from the line of nodes in the extruded direction. Since one-dimensional matrices are small, they are loaded into shared memory and reused for each triangular slice. Triangular matrices are then applied by loading matrix entries from global memory once and reusing them for each slice of the wedge. Algorithm 1 describes the implementation of the volume kernel in more detail.

For the surface kernels, each thread in a thread block is assigned to compute the right hand side contribution for a single node in the wedge. Since the lift matrices are Kronecker products of $\boldsymbol{L}^{\text{tri},k}$ and one-dimensional matrices or vectors, the computation is again broken up into an application of $\boldsymbol{L}^{\text{tri},k}$ and computations with one-dimensional matrices. First, numerical fluxes are computed and premultiplied by the triangular lift $\boldsymbol{L}^{\text{tri},k}$, which is loaded from global memory. These results are saved to shared memory, and used to compute the second half of the Kronecker product $\boldsymbol{L}^{\text{tri},k} \otimes \left(\widehat{\boldsymbol{M}}^{\text{1D}}\right)^{-1} \boldsymbol{e}$, which corresponds to scaling of the numerical fluxes (after premultiplication by the triangular lift matrix) by the jth entry of the vector $\left(\widehat{\boldsymbol{M}}^{\text{1D}}\right)^{-1} \boldsymbol{e}$ where j corresponds to the index of the current triangular slice of the wedge. The lift matrices for quadrilateral faces of the wedge are then directly applied to (non-premultiplied) numerical fluxes to complete the surface kernel. This procedure is outlined in Algorithm 2.

Since the loop sizes over which the surface kernel is parallelized vary between the number of total face nodes, the number of triangular nodes, and the number of nodes in the wedge, we take the number of threads to be the maximum of all these values. Consequentially, some threads remain idle at each iteration. However, experiments with parallelization over a smaller number of threads to reduce the number of idle threads (for example, over the number of face nodes in a triangular or quadrilateral face) resulted in slower overall runtimes.

Algorithm 1 Wedge volume kernel

parfor each element D^k **do**
  parfor triangular nodes $i = 1, \ldots, \frac{(N+1)(N+2)}{2}$ **do**
    for one dimensional nodes $j = 1, \ldots, N+1$ **do**
      Load geometric factors and solution variables from global memory into shared memory.
      Load one-dimensional matrix $\boldsymbol{D}_t^{\mathrm{1D}}$ to shared memory.
  end parfor
end parfor
Memory fence (synchronize threads).
parfor each element D^k **do**
  parfor triangular nodes $i = 1, \ldots, \frac{(N+1)(N+2)}{2}$ **do**
    Apply one-dimensional matrix $\boldsymbol{D}_t^{\mathrm{1D}}$ to the line of nodes extruded from node i.
    Store result to shared memory.
  end parfor
end parfor
Memory fence (synchronize threads).
parfor each element D^k **do**
  parfor triangular nodes $i = 1, \ldots, \frac{(N+1)(N+2)}{2}$ **do**
    Load triangular matrices $\boldsymbol{L}^{\mathrm{tri},k}, \boldsymbol{D}_r^{\mathrm{tri}}, \boldsymbol{D}_s^{\mathrm{tri}}$ from global memory.
    for one dimensional nodes $j = 1, \ldots, N+1$ **do**
      Apply triangular matrices to the jth triangular slice.
      Accumulate results into thread-local memory.
    for one dimensional nodes $j = 1, \ldots, N+1$ **do**
      Load results from thread-local memory.
      Assemble right hand side contributions and write to global memory.
  end parfor
end parfor

6.4. Computational evaluation and performance

In this section, we present results and timings which quantify the efficiency of a GPU-accelerated implementation of a nodal discontinuous Galerkin solver for vertically mapped wedge meshes. All computational results are reported for single precision. For the simulation of wave propagation in seismic applications, single precision is sufficient for both accuracy and numerical stability. For more complex nonlinear physics, some settings with highly disparate length scales benefit more from double precision. The use of double precision tends to result in slightly less than half the performance of GPU kernels in single precision, though the precise performance difference depends on the architecture of the specific GPU.

The computational complexity of both the volume and surface kernel for a given wedge element is $O(N^4)$. For the volume kernel, the cost is dominated by the application of the differentiation matrices. Since each differentiation matrix is a Kronecker product of a $\frac{(N+1)(N+2)}{2} \times \frac{(N+1)(N+2)}{2}$ triangular matrix and a one-dimensional $(N+1) \times (N+1)$ matrix, the cost is dominated by the $O(N^4)$ cost of applying triangular matrices. Similarly, for the surface kernel, the cost is dominated by the application of the lift matrices for each triangular face. These matrices are again representable as Kronecker products of triangular and one-dimensional matrices, implying that the asymptotic cost of the surface kernel is also $O(N^4)$.

For both the volume and surface kernels, computational performance comparable to standard volume and surface kernels for nodal discontinuous Galerkin methods on tetrahedra is observed [4, 36], though the achieved GFLOPS are slightly lower. In contrast, the observed bandwidth of the wedge volume and surface kernels is higher than that observed for nodal DG methods on tetrahedra by a factor of two for the volume kernel and a factor of slightly less than three for the surface kernel at $N = 5$.

Finally, comparing wall-clock runtimes of both wedge and tetrahedral kernels in Figure 11, we find that

Algorithm 2 Wedge surface kernel

parfor each element D^k **do**
  parfor nodes $i = 1, \ldots,$ total number of face nodes **do**
    Load geometric factors from global memory into shared memory.
    Load quadrilateral face lift matrices into shared memory.
  end parfor
end parfor
Memory fence (synchronize threads).
parfor each element D^k **do**
  parfor nodes $i = 1, \ldots,$ total number of face nodes **do**
    Load solution traces and compute numerical fluxes.
  end parfor
end parfor
Memory fence (synchronize threads).
parfor each element D^k **do**
  parfor triangular nodes $i = 1, \ldots, \frac{(N+1)(N+2)}{2}$ **do**
    Premultiply numerical fluxes on triangular faces by $J_f^k \boldsymbol{L}^{\mathrm{tri},k}$.
    Store result to shared memory.
  end parfor
end parfor
Memory fence (synchronize threads).
parfor each element D^k **do**
  parfor wedge nodes $i = 1, \ldots, N_p$ **do**
    Let j to be the index of the triangular slice node i lies in.
    Load premultiplied numerical fluxes on triangular faces, scale by jth entry of vector $\left(\widehat{\boldsymbol{M}}^{\mathrm{1D}}\right)^{-1} \boldsymbol{e}$.
    Apply lift matrices to numerical fluxes on quadrilateral faces.
    Assemble right hand side contributions and write to global memory.
  end parfor
end parfor

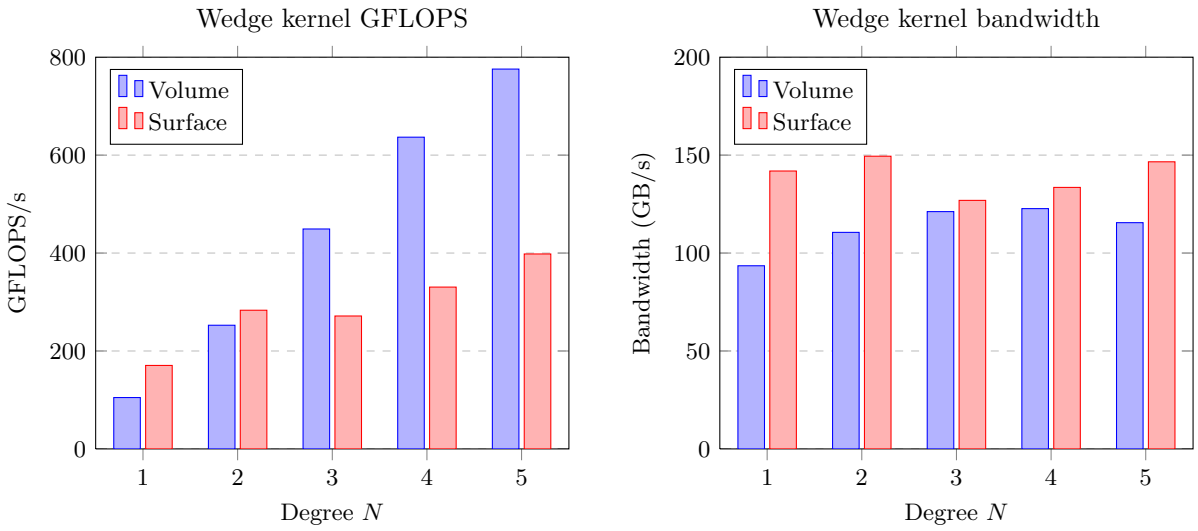

Figure 10: Computational performance of wedge kernels.

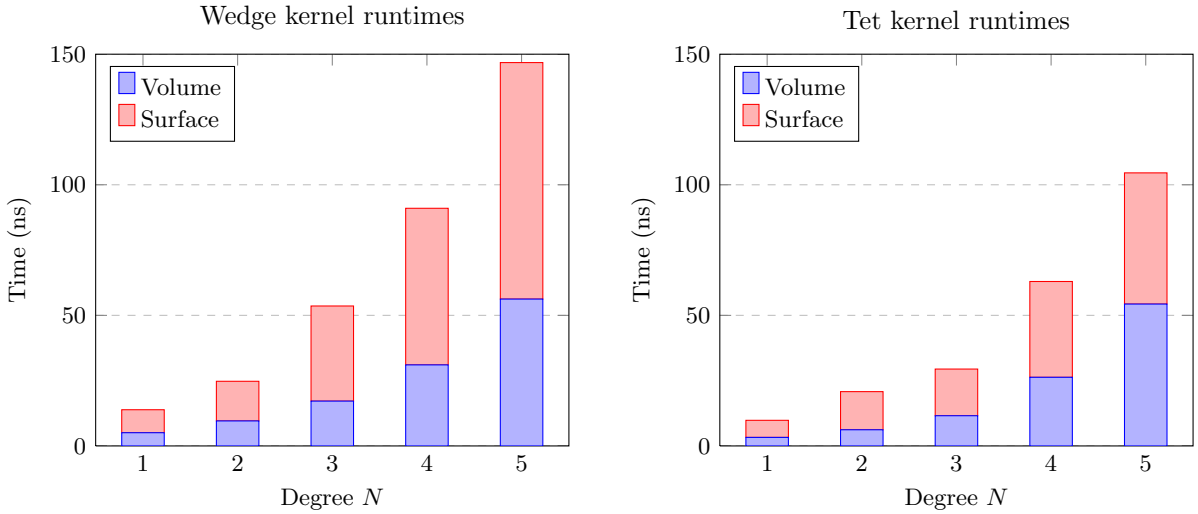

Figure 11: Per-element runtimes for wedge and tet kernels.

the per-element runtime of the wedge kernels is always greater than or equal to that of the tetrahedral kernels. However, this can be misleading, as the number of degrees of freedom for a high order wedge is greater than that of a high order tetrahedra. This is also reflected in the fact that a given volume can typically be triangulated with fewer wedges than tetrahedra; for example, a cube may be split into two well-conditioned prisms or five well-conditioned tetrahedra.

Normalizing by the number of degrees of freedom in Figure 12 shows that the per-dof runtime of the wedge volume kernel is lower than the per-dof runtime of the tetrahedral volume kernel for $N > 1$. For the surface kernel, the per-dof runtime is less for all orders except $N = 3$, for which tetrahedral kernels have been observed to perform better due to the number of degrees of freedom being close to a multiple of the number of concurrently executable threads in a workgroup [4].

We observe that the wedge volume kernel becomes more and more efficient relative to the tetrahedral volume kernel as N increases; this is due to the fact that the asymptotic complexity volume kernel is $O(N^4)$ for the wedge, as opposed to $O(N^6)$ for tetrahedra. We do not expect the same speedup as N increases for the surface kernel, since both the wedge and tetrahedral surface kernels have the same asymptotic cost of $O(N^4)$.

We conclude by presenting a computational simulation of wave propagation through the bi-unit cube $[-1, 1]^3$ with a wavespeed which is discontinuous across a complex interface. The domain is triangulated using a hybrid mesh consisting of vertically mapped wedges and unstructured tetrahedra, and the initial condition is set to be a Gaussian pulse located at the origin $(0, 0, 0)$. Figure 13 shows the pressure at time $t = .75$ computed using a DG method with order $N = 4$.

7. Conclusions and future work

Existing options for efficient low-storage discontinuous Galerkin methods on meshes with general mapped wedges are currently limited to Low-storage Curvilinear DG methods. Since LSC-DG methods are typically more expensive and less robust to mesh perturbations than quadrature-free nodal DG methods, we have developed in this work a limited-storage nodal DG method for vertically mapped wedges to address these issues. This method is also energy stable, in contrast to the mass-lumped methods typically employed for high order finite element methods on wedges. Additional interpolation to quadrature points is also avoided, despite the fact that exact quadrature is used. Finally, the method is shown to be computationally efficient,

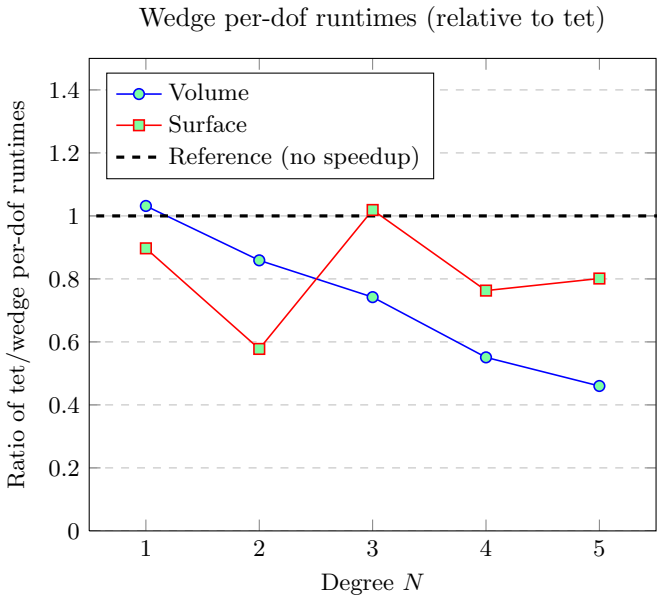

Figure 12: Ratio of per-dof runtimes of wedge kernels relative to per-dof runtimes of tetrahedral kernels (lower is better).

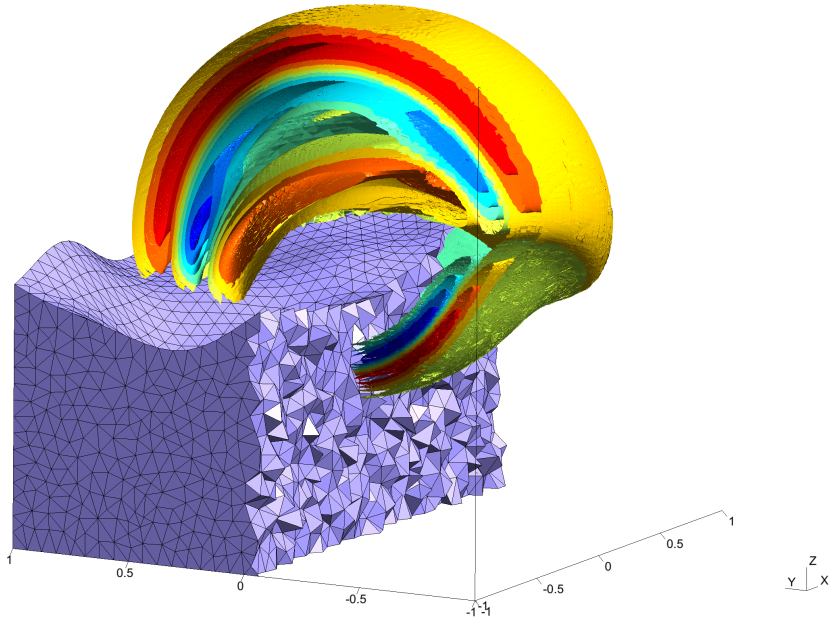

Figure 13: Numerical solution of the wave equation for a hybrid wedge-tet mesh ($N = 4$) with discontinuous wavespeed across a wavy interface.

with kernels attaining good computational performance and bandwidth at high orders of approximation, while reducing the per-dof runtime relative to tetrahedral kernels.

It is possible to achieve further speedup by taking advantage of the fact that the wedge meshes used in this work are structured in the vertical direction. This can be exploited to increase data locality by utilizing a layer-by-layer marching algorithm and reusing data from the previous layer [37, 38].

8. Acknowledgements

The authors thank TOTAL for permission to publish. JC and TW are funded by a grant from TOTAL E&P Research and Technology USA. Solutions were rendered using the high order visualization of GMSH [39].

References

[1] Dimitri Komatitsch and Jeroen Tromp. Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophysical journal international*, 139(3):806–822, 1999.

[2] George Karniadakis and Spencer Sherwin. *Spectral/hp element methods for computational fluid dynamics*. Oxford University Press, 2013.

[3] Jan S Hesthaven and Tim Warburton. *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*, volume 54. Springer, 2007.

[4] Andreas Klöckner, Tim Warburton, Jeff Bridge, and Jan S Hesthaven. Nodal discontinuous Galerkin methods on graphics processors. *Journal of Computational Physics*, 228(21):7863–7882, 2009.

[5] Nico Gödel, Nigel Nunn, Tim Warburton, and Markus Clemens. Scalability of higher-order discontinuous Galerkin FEM computations for solving electromagnetic wave propagation problems on GPU clusters. *Magnetics, IEEE Transactions on*, 46(8):3469–3472, 2010.

[6] A Modave, A St-Cyr, WA Mulder, and T Warburton. A nodal discontinuous galerkin method for reverse-time migration on gpu clusters. *Geophysical Journal International*, 203(2):1419–1435, 2015.

[7] Axel Modave, Amik St-Cyr, and Tim Warburton. Gpu performance analysis of a nodal discontinuous galerkin method for acoustic and elastic models. *Computers & Geosciences*, 91:64–76, 2016.

[8] Jesse Chan, Zheng Wang, Axel Modave, Jean-Francois Remacle, and T Warburton. GPU-accelerated discontinuous Galerkin methods on hybrid meshes. *Journal of Computational Physics*, 318:142–168, 2016.

[9] T. Warburton. A low-storage curvilinear discontinuous Galerkin method for wave problems. *SIAM Journal on Scientific Computing*, 35(4):A1987–A2012, 2013.

[10] Serguei Lapin, Sonja Kisin, Hua-wei Zhou, et al. Joint VSP and surface seismic tomography. In *2003 SEG Annual Meeting*. Society of Exploration Geophysicists, 2003.

[11] Qiang Wang, Sergey Danilov, and Jens Schröter. Finite element ocean circulation model based on triangular prismatic elements, with application in studying the effect of topography representation. *Journal of Geophysical Research: Oceans*, 113(C5), 2008.

[12] Sébastien Blaise, Richard Comblen, Vincent Legat, Jean-François Remacle, Eric Deleersnijder, and Jonathan Lambrechts. A discontinuous finite element baroclinic marine model on unstructured prismatic meshes. *Ocean Dynamics*, 60(6):1371–1393, 2010.

[13] Lorenzo Botti. Influence of reference-to-physical frame mappings on approximation properties of discontinuous piecewise polynomial spaces. *Journal of Scientific Computing*, 52(3):675–703, 2012.

[14] T Warburton. An explicit construction of interpolation nodes on the simplex. *Journal of engineering mathematics*, 56(3):247–262, 2006.

[15] RM Kirby, TC Warburton, I Lomtev, and GE Karniadakis. A discontinuous Galerkin spectral/hp method on hybrid grids. *Applied numerical mathematics*, 33(1):393–405, 2000.

[16] Morgane Bergot and Marc Duruflé. Higher-order discontinuous Galerkin method for pyramidal elements using orthogonal bases. *Numerical Methods for Partial Differential Equations*, 29(1):144–169, 2013.

[17] David A Kopriva and Gregor Gassner. On the quadrature and weak form choices in collocation type discontinuous Galerkin spectral element methods. *Journal of Scientific Computing*, 44(2):136–155, 2010.

[18] Jesse Chan, Russell J Hewett, and T Warburton. Weight-adjusted discontinuous Galerkin methods: wave propagation in heterogeneous media. *arXiv preprint arXiv:1608.01944*, 2016.

[19] Jan S Hesthaven and T. Warburton. Nodal high-order methods on unstructured grids: I. time-domain solution of Maxwell's equations. *Journal of Computational Physics*, 181(1):186–221, 2002.

[20] Jan S Hesthaven and T. Warburton. High-order nodal discontinuous Galerkin methods for the Maxwell eigenvalue problem. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 362(1816):493–524, 2004.

[21] Gregor J Gassner. A skew-symmetric discontinuous Galerkin spectral element discretization and its relation to SBP-SAT finite difference methods. *SIAM Journal on Scientific Computing*, 35(3):A1233–A1253, 2013.

[22] Rajesh Gandham. *High performance high-order numerical methods: applications in ocean modeling*. PhD thesis, Rice University, 2015.

[23] Susanne Brenner and Ridgway Scott. *The mathematical theory of finite element methods*, volume 15. Springer Science & Business Media, 2007.

[24] Jesse Chan and T Warburton. Orthogonal bases for vertex-mapped pyramids. *SIAM Journal on Scientific Computing*, 38(2):A1146–A1170, 2016.

[25] Hui Wan, Marco A Giorgetta, Günther Zängl, Marco Restelli, Detlev Majewski, Luca Bonaventura, Kristina Fröhlich, Daniel Reinert, P Rıpodas, Luis Kornblueh, et al. The ICON-1.2 hydrostatic atmospheric dynamical core on triangular grids, Part I: formulation and performance of the baseline version. *Geoscientific Model Development*, 6:735–763, 2013.

[26] Vincenzo Casulli and Roy A. Walters. An unstructured grid, three-dimensional model based on the shallow water equations. *International journal for numerical methods in fluids*, 32(3):331–348, 2000.

[27] Norman A Phillips. A coordinate system having some special advantages for numerical forecasting. *Journal of Meteorology*, 14(2):184–185, 1957.

[28] Hang Si. TetGen, a Delaunay-based quality tetrahedral mesh generator. *ACM Transactions on Mathematical Software (TOMS)*, 41(2):11, 2015.

[29] T Warburton. A low storage curvilinear discontinuous Galerkin time-domain method for electromagnetics. In *Electromagnetic Theory (EMTS), 2010 URSI International Symposium on*, pages 996–999. IEEE, 2010.

[30] David S Medina, Amik St-Cyr, and T Warburton. OCCA: A unified approach to multi-threading languages. *arXiv preprint arXiv:1403.0968*, 2014.

[31] Bernardo Cockburn, Bo Dong, and Johnny Guzmán. Optimal convergence of the original DG method for the transport-reaction equation on special meshes. *SIAM Journal on Numerical Analysis*, 46(3):1250–1265, 2008.

[32] Claes Johnson and Juhani Pitkäranta. An analysis of the discontinuous Galerkin method for a scalar hyperbolic equation. *Mathematics of computation*, 46(173):1–26, 1986.

[33] Lucas C Wilcox, Georg Stadler, Carsten Burstedde, and Omar Ghattas. A high-order discontinuous Galerkin method for wave propagation through coupled elastic–acoustic media. *Journal of Computational Physics*, 229(24):9373–9396, 2010.

[34] Douglas Arnold, Daniele Boffi, and Richard Falk. Approximation by quadrilateral finite elements. *Mathematics of computation*, 71(239):909–922, 2002.

[35] Nico Gödel, Steffen Schomann, Tim Warburton, and Markus Clemens. GPU accelerated Adams–Bashforth multirate discontinuous Galerkin FEM simulation of high-frequency electromagnetic fields. *Magnetics, IEEE Transactions on*, 46(8):2735–2738, 2010.

[36] Jesse Chan and T Warburton. GPU-accelerated Bernstein-Bezier discontinuous Galerkin methods for wave problems. *arXiv preprint arXiv:1512.06025*, 2015.

[37] Paulius Micikevicius. 3D finite difference computation on GPUs using CUDA. In *Proceedings of 2nd workshop on general purpose processing on graphics processing units*, pages 79–84. ACM, 2009.

[38] Gheorghe-Teodor Bercea, Andrew T. T. McRae, David A. Ham, Lawrence Mitchell, Florian Rathgeber, Luigi Nardi, Fabio Luporini, and Paul H. J. Kelly. A numbering algorithm for finite element on extruded meshes which avoids the unstructured mesh penalty. *arXiv preprint arXiv:1604.05937*, 2015.

[39] Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009.